arXiv:1611.08022v1 [math.OC] 23 Nov 2016

A STRONGER CONVERGENCE RESULT ON THE PROXIMAL INCREMENTAL AGGREGATED GRADIENT METHOD

N. D. VANLI*, M. GÜRBÜZBALABAN†, AND A. OZDAGLAR†

Abstract. We study the convergence rate of the proximal incremental aggregated gradient (PIAG) method for minimizing the sum of a large number of smooth component functions (where the sum is strongly convex) and a non-smooth convex function. At each iteration, the PIAG method moves along an aggregated gradient formed by incrementally updating gradients of component functions at least once in the last K iterations and takes a proximal step with respect to the non-smooth function. We show that the PIAG algorithm attains an iteration complexity that grows linear in the condition number of the problem and the delay parameter K. This improves upon the previously best known global linear convergence rate of the PIAG algorithm in the literature which has a quadratic dependence on K.

1. Introduction. We consider composite additive cost optimization problems, where the objective function is given by the sum of m component functions $f_i(x)$ and a possibly non-smooth regularization function $r(x)$:

$$\min_{x\in\mathbb{R}^n} F(x) \triangleq f(x) + r(x), \tag{1.1}$$

where $f(x) = \frac{1}{m}\sum_{i=1}^m f_i(x)$. Each component function $f_i : \mathbb{R}^n \to (-\infty, \infty)$ is assumed to be convex and continuously differentiable while the regularization function $r : \mathbb{R}^n \to (-\infty, \infty]$ is proper, closed, and convex but not necessarily differentiable.

The recent paper [4] studied the PIAG algorithm, which at each iteration $k \geq 0$, first constructs an *aggregated gradient* defined by

$$g_k \triangleq \frac{1}{m}\sum_{i=1}^m \nabla f_i(x_{\tau_{i,k}}),$$

where $\nabla f_i(x_{\tau_{i,k}})$ represents the gradient of the ith component function sampled at time $\tau_{i,k}$. This aggregated gradient is used to update x_k as

$$x_{k+1} = \text{prox}_r^\eta(x_k - \eta g_k), \tag{1.2}$$

where the proximal mapping is defined as $\text{prox}_r^\eta(y) = \arg\min_{x\in\mathbb{R}^n}\left\{\frac{1}{2}||x-y||^2 + \eta r(x)\right\}$ with a constant step size $\eta > 0$.

It was shown in [4] that the PIAG algorithm attains a global linear convergence rate of $1 - \widetilde{\mathcal{O}}(Q^{-1}K^{-2})$ in function suboptimality, i.e., $F(x_k) - F(x^*)$, where x^* denotes the optimal solution to (1.1) and the tilde is used to hide the logarithmic terms in Q and K. This result implies that in order

*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. email: {denizcan, mertg, asuman}@mit.edu.

to achieve an ϵ-optimal solution, PIAG requires at most $\widetilde{\mathcal{O}}(QK^2 \log(1/\epsilon))$ iterations. The independent work [1] also studied the PIAG algorithm and showed using a different analysis that it attains a global linear convergence rate of $1 - \mathcal{O}(Q^{-1}K^{-2})$ in distance to the optimal solution $||x_k - x^*||$. This result implies that to achieve a point in the ϵ-neighborhood of the optimal solution, PIAG requires $\mathcal{O}(QK^2 \log(1/\epsilon))$ iterations. The latter result on distances does not translate directly into a linear convergence rate in function suboptimality since the problem (1.1) is not smooth.

In this paper, by using the results presented in [1] and [4], we provide a stronger linear convergence rate for the deterministic PIAG algorithm. In particular, in [4], two lemmas regarding the relations on the proximal operator are introduced to provide a contraction relation on the function suboptimality. In [1], a lemma that characterizes the linear convergence of a Lyapunov function is introduced, where the Lyapunov function satisfies a certain contraction relation with sufficiently small perturbation that depends on the recent history. By using these results, we prove that the PIAG algorithm attains a global linear convergence rate of $1 - \mathcal{O}(Q^{-1}K^{-1})$ in function suboptimality. This implies that in order to achieve an ϵ-optimal solution in suboptimality of the function values, PIAG requires at most $\mathcal{O}(QK \log(1/\epsilon))$ iterations. To our knowledge, this convergence rate result provides the best dependence on the condition number of the problem Q and the delay parameter K for deterministic incremental aggregated gradient methods.

2. Assumptions. Throughout the paper, we make the following standard assumptions that are used in both [1] and [4].

ASSUMPTION 2.1. ***(Lipschitz gradients)*** *Each f_i has Lipschitz continuous gradients on $\mathbb{R}^n$ with some constant $L_i \geq 0$, i.e.,*

$$||\nabla f_i(x) - \nabla f_i(y)|| \leq L_i \, ||x - y|| \,,$$

for any $x, y \in \mathbb{R}^n$.[1]

Defining $L \triangleq \frac{1}{m} \sum_{i=1}^m L_i$, we observe that Assumption 2.1 and the triangle inequality yield

$$||\nabla f(x) - \nabla f(y)|| \leq L \, ||x - y|| \,,$$

for any $x, y \in \mathbb{R}^n$, i.e., the function f is L-smooth.

ASSUMPTION 2.2. ***(Strong Convexity)*** *The sum function f is μ-strongly convex on $\mathbb{R}^n$ for some $\mu > 0$, i.e., the function $x \mapsto f(x) - \frac{\mu}{2} ||x||^2$ is convex.*

ASSUMPTION 2.3. ***(Subdifferentiability)*** *The regularization function $r : \mathbb{R}^n \to (-\infty, \infty]$ is proper, closed, convex and subdifferentiable everywhere in its effective domain, i.e., $\partial r(x) \neq \emptyset$ for all $x \in \{y \in \mathbb{R}^n : r(y) < \infty\}$.*

[1] If a function f has Lipschitz continuous gradients with some constant L, then f is called L-smooth. We use these terms interchangeably.

A consequence of Assumptions 2.2 and 2.3 is that F is strongly convex, hence there exists a unique optimal solution of problem (1.1), which we denote by x^* (cf. Lemma 6 in [3]).

Another consequence of Assumption 2.3 is that the set of subgradients of x_k is well-defined for all $k \geq 0$. Then, it follows from the optimality conditions [2] of the minimization problem in the proximal map in (1.2) that $0 \in \partial\phi(x_{k+1})$. This yields $x_{k+1} - (x_k - \eta g_k) + \eta h_{k+1} = 0$, for some subgradient $h_{k+1} \in \partial r(x_{k+1})$. Thus, we can represent the update rule of the PIAG algorithm as follows

$$x_{k+1} = x_k + \eta d_k,$$

where $d_k \triangleq -g_k - h_{k+1}$ is the direction of the update at time k.

ASSUMPTION 2.4. ***(Bounded Delay)*** *Each component function is sampled at least once in the past $K \geq 0$ iterations, i.e., there exists a finite integer K such that $k - K \leq \tau_{i,k} \leq k$, for all $k \geq 1$ and $i \in \{1, \ldots, m\}$.*

3. Main Result. In this section, we characterize the global linear convergence rate of the PIAG algorithm. Let

$$F_k \triangleq F(x_k) - F(x^*) \tag{3.1}$$

denote the suboptimality in the objective value at iteration k. The paper [4] presented two lemmas regarding the evolution of F_k and $||d_k||^2$. In particular, the first lemma investigates how the suboptimality in the objective value evolves over the iterations and the second lemma relates the direction of update to the suboptimality in the objective value at a given iteration k.

LEMMA 3.1. *[4, Lemma 3.3] Suppose that Assumptions 2.1-2.4 hold. Then, the PIAG algorithm yields the following guarantee*

$$F_{k+1} \leq F_k - \frac{1}{2}\eta\,||d_k||^2 + \eta^2 \frac{L}{2} \sum_{j=(k-K)_+}^{k-1} ||d_j||^2\,, \tag{3.2}$$

for any step size $0 < \eta \leq \frac{1}{L(K+1)}$.

LEMMA 3.2. *[4, Lemma 3.5] Suppose that Assumptions 2.1-2.4 hold. Then, for any $0 < \eta \leq \frac{1}{L(K+1)}$, the PIAG algorithm yields the following guarantee*

$$-||d_k||^2 \leq -\frac{\mu}{4} F_{k+1} + \eta L \sum_{j=(k-K)_+}^{k-1} ||d_j||^2\,.$$

Before presenting the main result of this work, we introduce the following lemma, which was presented in [1], in a slightly different form. This lemma shows linear convergence rate for a nonnegative

sequence Z_k that satisfies a contraction relation perturbed by shocks (represented by Y_k in the lemma).

LEMMA 3.3. *[1, Lemma 1] Let $\{Z_k\}$ and $\{Y_k\}$ be a sequence of non-negative real numbers satisfying*

$$\alpha Z_{k+1} \leq Z_k - \beta Y_k + \gamma \sum_{j=k-A}^{k} Y_j, \tag{3.3}$$

for any $k \geq 0$ for some constants $\alpha > 1$, $\beta \geq 0$, $\gamma \geq 0$ and $A \in \mathbb{Z}^+$. If

$$\gamma(\alpha^{A+1} - 1) \leq \beta(\alpha - 1) \tag{3.4}$$

holds, then $Z_k \leq \alpha^{-k} Z_0$, for all $k \geq 0$.

We next present the main theorem of this paper, which characterizes the linear convergence rate of the PIAG algorithm. In particular, we show that when the step size is sufficiently small, the PIAG algorithm is linearly convergent with a contraction rate that depends on the step size η and the strong convexity constant μ.

THEOREM 3.4. *Suppose that Assumptions 2.1-2.4 hold. Then, the PIAG algorithm with step size $0 < \eta \leq \frac{16}{\mu}\left[\left(1 + \frac{1}{48Q}\right)^{\frac{1}{K+1}} - 1\right]$ is linearly convergent satisfying*

$$F_k \leq \left(1 + \eta\frac{\mu}{16}\right)^{-k} F_0, \tag{3.5}$$

for any $k \geq 0$. Furthermore, if $\eta = \frac{16}{\mu}\left[\left(1 + \frac{1}{48Q}\right)^{\frac{1}{K+1}} - 1\right]$, then

$$F_k \leq \left(1 - \frac{1}{49Q(K+1)}\right)^{k} F_0. \tag{3.6}$$

Proof. By Lemma 3.2, we have

$$-\frac{1}{4}\eta \left|\left|d_k\right|\right|^2 \leq -\eta\frac{\mu}{16}F_{k+1} + \eta^2\frac{L}{4}\sum_{j=(k-K)_+}^{k-1} \left|\left|d_j\right|\right|^2.$$

Using this inequality in (3.2) of Lemma 3.1, we get

$$\left(1 + \eta\frac{\mu}{16}\right) F_{k+1} \leq F_k - \frac{1}{4}\eta \left|\left|d_k\right|\right|^2 + \eta^2\frac{3L}{4}\sum_{j=(k-K)_+}^{k-1} \left|\left|d_j\right|\right|^2. \tag{3.7}$$

We want to apply Lemma 3.3 to the inequality (3.7) with $Z_k = F_k$ and $Y_k = \left|\left|d_k\right|\right|^2$ for proving (3.5). For this purpose, we need $0 < \eta \leq \frac{1}{L(K+1)}$ in order for Lemma 3.1 and Lemma 3.2 to hold, and

$$\eta^2\frac{3L}{4}\left(\left(1 + \eta\frac{\mu}{16}\right)^{K+1} - 1\right) \leq \frac{1}{4}\eta\left(\left(1 + \eta\frac{\mu}{16}\right) - 1\right) \tag{3.8}$$

with $\eta > 0$ for Lemma 3.3 to hold. Simplifying and rearranging terms in (3.8), we obtain

$$\left(1 + \eta\frac{\mu}{16}\right)^{K+1} - 1 \leq \frac{1}{48Q}.$$

Therefore, for any step size satisfying

$$0 < \eta \leq \frac{16}{\mu}\left[\left(1+\frac{1}{48Q}\right)^{\frac{1}{K+1}} - 1\right], \tag{3.9}$$

Lemma 3.3 holds. We can also observe that the right-hand side of (3.9) can be upper bounded using the Bernoulli inequality, i.e., $(1+x)^r \leq 1+rx$ for any $x \geq -1$ and $r \in [0,1]$, as follows

$$\begin{aligned}\eta &\leq \frac{16}{\mu}\left(1+\frac{1}{48Q(K+1)} - 1\right)\\ &= \frac{1}{3L(K+1)}.\end{aligned} \tag{3.10}$$

Thus, the constraint (3.9) satisfies the constraint $0 < \eta \leq \frac{1}{L(K+1)}$ in Lemma 3.1 and Lemma 3.2 as well. Then, applying Lemma 3.3 to (3.7) yields (3.5).

It remains to show (3.6). Plugging in $\eta = \frac{16}{\mu}\left[\left(1+\frac{1}{48Q}\right)^{\frac{1}{K+1}} - 1\right]$ in (3.5) of Theorem 3.4, we have

$$\begin{aligned}F_k &\leq \left(1+\frac{16}{\mu}\left[\left(1+\frac{1}{48Q}\right)^{\frac{1}{K+1}} - 1\right]\frac{\mu}{16}\right)^{-k} F_0\\ &= \left(1+\frac{1}{48Q}\right)^{\frac{-k}{K+1}} F_0\\ &\leq \left(1-\frac{1}{49Q}\right)^{\frac{k}{K+1}} F_0\\ &\leq \left(1-\frac{1}{49Q(K+1)}\right)^{k} F_0,\end{aligned} \tag{3.11}$$

where the third line follows as $Q \geq 1$. This implies (3.6) and completes the proof. □

We next introduce the following corollary, which highlights the main result of the paper. This corollary indicates that for an appropriately chosen step size, the PIAG algorithm is guaranteed to return an ϵ-optimal solution after $\mathcal{O}(QK\log(1/\epsilon))$ iterations.

COROLLARY 3.5. *Suppose that Assumptions 2.1-2.4 hold. Then, the PIAG algorithm with step size* $\eta = \frac{16}{\mu}\left[\left(1+\frac{1}{48Q}\right)^{\frac{1}{K+1}} - 1\right]$ *is guaranteed to return an* ϵ*-optimal solution after at most* $49Q(K+1)\log(F_0/\epsilon)$ *iterations.*

Proof. By Theorem 3.4, the inequality (3.6) holds. Taking logarithm of both sides of this inequality yields

$$\begin{aligned}\log(F_k) &\leq \log(F_0) + k\log\left(1-\frac{1}{49Q(K+1)}\right)\\ &\leq \log(F_0) - \frac{k}{49Q(K+1)},\end{aligned}$$

where the last line follows since $\log(1+x) \leq x$ for any $x \geq -1$. Therefore, for any k satisfying

$$\log(F_0) - \frac{k}{49Q(K+1)} \leq \log(\epsilon), \tag{3.12}$$

x_k is an ϵ-optimal solution. Rearranging terms in (3.12), we conclude that for any $k \geq 49Q(K+1)\log(F_0/\epsilon)$, x_k is an ϵ-optimal solution. □

REFERENCES

[1] A. Aytekin, H. R. Feyzmahdavian, and M. Johansson. Analysis and implementation of an asynchronous optimization algorithm for the parameter server. *CoRR*, abs/1610.05507, 2016.

[2] D. P. Bertsekas. *Convex Optimization Algorithms*. Athena Scientific, 2015.

[3] Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, 2009.

[4] N. D. Vanli, M. Gurbuzbalaban, and A. Ozdaglar. Global convergence rate of proximal incremental aggregated gradient methods. *CoRR*, abs/1608.01713, 2016.